

AD-A269 054



Q
100

Gigabit Network Communications Research

Quarterly R&D Status Report No. 10

Period: 1 JAN 93- 31 MAR 93

DTIC QUALITY INSPECTED 1

Accession For	
NTIS	CRA&I
DTIC	TAB
Unannounced	
Justification: <i>#A269054</i>	
By _____	
Distribution / _____	
Availability Codes	
Dist	Avail and/or Special
A-1	

DTIC
ELECTED
SEP 01 1993
S E D

Sponsored by
Advanced Research Projects Agency (DoD)
Computer Systems Technology Office
Gigabit Network Communications Research
Under Contract #DABT63-91-C-0001
PR&C: HR0011-0218-0001
AAP No. DAR1010
Issued by Directorate of Contracting
Fort Huachuca, AZ

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U.S. Government.

Approved for public release
Distribution Unlimited

93 8 24 119

93-19790
Bp8

GIGABIT NETWORK COMMUNICATIONS RESEARCH

QUARTERLY R&D STATUS REPORT for DABT63-91-C-0001

(covering the period 1/1/93-3/31/93)

1. ATOMIC

The ATOMIC project demonstrates the viability of utilizing inexpensive message-passing, fine grain, supercomputer components to create a very high-speed local area network (VHSLAN).

1.A. SOFTWARE

Support for Extended-Windows TCP was added to ATOMIC and tested in conjunction with the NSF project on Protocol Engineering.

Behavior of device driver timings was studied. The interrupt priority of some of the SBus boards was adjusted and tests were performed. The results were inconclusive.

To increase reliability of the network, an ATOMIC packet header is now used on all packets. This includes a checksum over the ATOMIC header for error control. In addition a UDP checksum facility was added to Mosaic code transfers across the network as an insurance check for errors during network bootstrapping.

Improvements were made to the ATOMIC multicasting code to enhance its reliability.

Preliminary investigations were undertaken to determine ways in which ATM could be integrated into ATOMIC.

1.B. HARDWARE

The prototype for an ATOMIC active cable board was designed and submitted for fabrication. The goal of this effort is to demonstrate an inexpensive 500 Mb/s twisted-pair ATOMIC cable that is capable of 30 meter runs. Board delivery is to occur in mid-April.

Since this design is for prototyping and testing, attention was paid to versatility. The design will allow insertion of either SLACK or Dialog chips. It will also allow testing the cable driver logic without either SLACK or Dialog installed. Characterization and testing will occur next quarter.

Operational SLACK chips arrived from Caltech at the end of March. These will be used in the prototype cable testing.

2. Personal Conferencing

Tools for teleconferencing and telecollaboration must support different requirements for a variety of session modes across a wide variation in session size, population size, and geographic distribution. There has been significant progress over the past year in experimentation at the large end of these scales where only a "loose control" style of conferencing is feasible. A collection of Internet packet audio and video tools from several institutions have been used for audio/video multicasts of large "seminar" events such as IETF meetings.

However, it is also important to address the needs of smaller conferences with an explicit set of participants for which a "tight control" style of session management is practical. Through explicit exchange of state information among the participants, several useful and necessary services can be added, including:

- an explicit request for participation by a callee, perhaps with a small pop-up window, and an explicit reply (no answer, busy/refuse, accept);
- key exchange to implement security measures such as authentication of conferees and data confidentiality and integrity;
- negotiation of session parameters, such as selecting a compatible compression algorithm and data rate.

To demonstrate these features using the same audio and video tools with both loose- and tight-control session management, we have modified our conference control program, MMCC, to invoke Xerox PARC's *uv* and LBL's *vat*, in addition to our own VT and PVP tools that we have utilized previously for experiments using both ST-II and IP multicast. We have also completed the conversion of MMCC to run under the X Window System and the XView library, to be compatible with the environment of these tools.

In a demonstration of this capability at the March IETF meeting, we used MMCC to establish conferences among a specific set of participants (vs. having a wide-open session). MMCC initiated sessions with a user-selected subset of the underlying tools, allocating multicast addresses and then distributing those to the selected set of participants along with a shared DES encryption key to implement confidentiality of the data stream. Since MMCC stays "in the loop" after session initiation, we expect that in the future it will provide a channel for negotiation of quality of service and configuration parameters as they may change during the session. At the end of the session, MMCC coordinated the shutdown of the underlying tools. We intend soon to release MMCC for more widespread use in the Internet as an alternative to the LBL *sd* session directory tool that implements loose-control session management.

In this first step for demonstration purposes, the session encryption key is exchanged in the clear. Our next steps will be to protect the session key with a manually prearranged private key, and then for the longer term we are investigating how to integrate mechanisms such as those in PEM or Kerberos with MMCC's reliable control communication paths. We are working to identify the range of security threats that may

exist for applications such as conferencing and to prioritize the utility and implementability of security services to combat them. Implementing these services will require a more sophisticated interface between session management and the audio/video tools than simple program initiation; we are also studying these interface requirements.

This quarter, we extended the teleconferencing facilities at ARPA and ISI by interconnecting the wide-area packet teleconferencing systems, which use DARTnet and DSInet, with the ZAPT local-area desktop conferencing system. The ZAPT system, installed both at ARPA and ISI, uses analog audio and video with NeXT workstations running a custom extension of Bellcore's Touring Machine software, and provides a local analog distribution and teleconferencing capability.

Using MMCC, conference users can select among different codecs for different wide area systems, such as DARTnet and DSInet, and among dedicated conference rooms or the interconnect to the desktop system. To avoid conflicts created by multiple conferencing systems needing access to shared hardware (e.g., both MMCC and ZAPT controlling the crossbar switch, echo canceller, video codecs), we integrated resource registration into our device servers. Now, a client application can use the server to reserve access to the hardware and prevent another teleconferencing application from stealing it. The resource reservation is also fault tolerant, so that owner failure releases the device.

At the March IETF meeting, we chaired two teleconferencing-related groups. The Audio/Video Transport Working Group is nearing completion in the specification of an experimental real-time transport protocol (RTP). We have implemented this draft protocol in the PVP video program, and we have implemented a routine for software decoding of the video data stream produced by the Bolter/Concept codec and PVP which was then integrated into PARC's *nv* video tool. Together, these developments allow the PVP and *nv* programs to interoperate, which makes it now possible for all DARTnet member sites to view each other with video in teleconferences.

The second IETF group is on Conferencing Control, and met as a BOF to discuss the requirements of a session protocol for multiway collaborations, and particularly to focus on a suitable scope for this problem to establish a charter and a working group. The intent is to specify a protocol to support loose- and tight-control conference styles.

This was the fourth live multicast of audio and video from an IETF meeting.

3. Integrated Services Protocols

The problem of Integrated Service (IS) in the Internet can be separated into a number of components, in particular: traffic control, admission control, and the setup protocol. ISI proposes for focus on "Reservation Protocol (RSVP)", especially on the protocol engineering of the setup protocol, in collaboration with Xerox PARC. In future reports, progress on Internet IS will be reported under "RSVP".

We have started to collaborate with Lixia Zhang of Xerox PARC to prepare a protocol specification document for RSVP. Under the RSVP task, we will also take responsibility for integrating a traffic control mechanism, CSZ, developed by MIT and Xerox PARC and an admission control algorithm developed by Xerox PARC, MIT, and USC, with RSVP, and carry out a DARTnet demonstration of IS supporting teleconferencing.

4. IP/SQ Congestion Control

Investigation of source quench continued this quarter. Results are expected next quarter. A technical report will be written.

5. Tunnel

The ISI Tunnel System is an access control gateway that performs packet filtering. The Tunnel provides smooth seamless initiation of communication from inside the gateway, while performing explicit access control on communication initiated outside of the gateway. The Tunnel System consists of a modified kernel and a set of utility programs, and is implemented on a SPARCstation, under Sun OS 4.1.2. The gateway portion of the Tunnel is implemented as an extension to the SunOS kernel, and performs IP-level filtering. The Tunnel System also includes a set of utility programs, which are used to set up access controls in the Tunnel, via system calls, and to monitor the status of the Tunnel. Changes in access control parameters and counts of packets forwarded and packets dropped, as well as any error conditions, are logged using the standard Unix *syslogd* facility.

This quarter we completed the logging functionality, as well as the utility programs that are required for the operation and the management of the Tunnel. In addition, we added "keepalive-avoidance" logic to the Tunnel's filtering mechanism to prevent users/applications from inadvertently prolonging automatic access through the Tunnel, through the use of keepalives. A limited port/protocol filtering capability has also been included to allow the use of a specific application, such as telnet, from outside of the gateway, to be disallowed.

Next quarter we will complete the testing of the Tunnel at ISI, in preparation for the planned installation of the Tunnel at ARPA.

6. Zoned Analog Personal Teleconferencing (ZAPT) formerly Automated Cluster Teleconferencing (ACT)

The ACT project has been renamed ZAPT (Zoned Analog Personal Teleconferencing), to avoid overlap with a satellite research project of a similar name (ACTS).

The known hardware and software faults have been identified and corrected. New switch control software now performs verify-writes, and logs errors in more detail. The

equipment at ARPA is operational, and has been installed for a total of 10 offices (of 15 available) and 1 external connection. The ISI-developed bridging extension has been verified, and has been delivered back to Bellcore.

ISI has also developed a "radio" user client to permit users to join a session in progress. This is an extension to the Touring Machine (TM) Version 2 functionality. The radio client is being expanded to permit automatic setup of non-Touring Machine components.

The ZAPT system is now more robust, due to the addition of a monitor called SafeNet. SafeNet starts the system components, and if any fail, automatically kills the remaining components and restarts the entire set and resets the hardware components. A NeXT GUI has been created for single-keyclick startup of the system, rather than command-line shell startup. The NeXT GUI incorporates a user version of SafeNet, so that if any components fail, the dependent components are killed, and the suite restarts. A GUI interface provides forced reset and quit (exit) capability.

Currently, we are reviewing the functionality of the system to determine the next phases of teleconferencing at ARPA. A document is under preparation describing ISI's contribution to the extension of TM. We are continuing the development of a multipoint multimedia model (M3) that integrates Bellcore's TM and ISI's MMCC models.

7. DARTnet Network Operations Center (DARTNOC)

DARTNOC continues to monitor and maintain DARTnet. Test times for experimenters are continuing to be scheduled. No Sprint related outages have occurred since January.

8. Infrastructure

8.A. USER SERVICES

As the User Services Area Director, Joyce Reynolds participated in IESG Teleconferences from January – March 1993 and attended the IETF meeting in Columbus, Ohio, March 29 – April 2, 1993.

Ten working groups in the User Services Area of the IETF met in Columbus, Ohio. Three BOFs (Birds of a Feather) were held in the User Services area. These included: World Wide Web, Gopher, and the Low-end IP Hardware Wish List.

During this period, one new Working Group was formed in the User Services Area of the IETF: Integrated Directory Services (IDS).

There are currently 11 active Working Groups in the User Services Area of the IETF.

Two papers were published as FYI RFCs:

FYI 18 "Internet Users' Glossary", (Also RFC 1392), January 1993.

FYI 17 "The Tao of IETF - A Guide for New Attendees of the Internet Engineering Task Force, (Also RFC 1391), January 1993.

8.B. INTERNET MONTHLY REPORT

The Internet Monthly Report (IMR) is the status report on the operation of the Internet and the research and development activities of the Internet community. It features reports from the IAB, the Internet Research Task Force and its research groups, and the Internet Engineering Task Force and its working groups in addition to the reports from approximately 30 regional networks and individual sites. A typical monthly report is approximately 40 pages.

During this reporting period, three Internet Monthly Reports for December 1992, January 1993, and February 1993 were assembled, edited, and distributed directly (via electronic mail) to over 375 mailboxes, some of which are exploder mailboxes where the report is sent to a sublist of people. In particular, the mailbox "IETF@isi.edu", which is one of the mailboxes on the IMR list, goes out to an additional 1,000 mailboxes, many of which are further exploders.

8.C. HIGH PERFORMANCE NETWORK RESEARCH REPORT

The High Performance Network Research Report (HPNRR) discusses research and development activities in the Gigabit Networking Research program and the advanced networking research community. A typical report is about 25 pages. During this reporting period, three reports for December 1992, January 1993, and February 1993 were assembled, edited, and distributed directly (via electronic mail) to over 130 people.

8.D. REQUEST FOR COMMENTS

ISI serves as the technical editor and "publisher" of the Internet document series called "Requests for Comments" (RFCs). 51 RFCs were published this quarter:

- RFC 1384: Barker, P., and S.E. Hardcastle-Kille, "Naming Guidelines for Directory Pilots", University College London, ISODE Consortium, February 1993.
- RFC 1387: Malkin, G., "RIP Version 2 Protocol Analysis", Xylogics Inc., January 1993.
- RFC 1388: Malkin, G., "RIP Version 2 Carrying Additional Information", Xylogics Inc., January 1993.
- RFC 1389: Malkin, G. (Xylogics Inc.), and F. Baker (Advanced Computer Communications), "RIP Version 2 MIB Extensions", January 1993.

- RFC 1390: Katz, D., "Transmission of IP and ARP over FDDI Networks", Cisco Systems, Inc., January 1993.
- RFC 1391: Malkin, G., "The TAO of IETF – A Guide for New Attendees of the Internet Engineering Task Force", Xylogics Inc., January 1993.
- RFC 1392: Malkin, G. (Xylogics Inc.), and T. LaQuey Parker (UTexas), "Internet Users' Glossary", January 1993.
- RFC 1393: Malkin, G., "Traceroute Using an IP Option", Xylogics Inc., January 1993.
- RFC 1394: Robinson, P., "Relationship of Telex Answerback Codes to Internet Domains", Tansin A. Darcos and Co., January 1993.
- RFC 1395: Reynolds, J.K., "BOOTP Vendor Information Extensions", USC/ISI, January 1993.
- RFC 1396: Crocker, S., "The Process for Organization of Internet Standards Working Group (POISED)", Trusted Information Systems, Inc., January 1993.
- RFC 1397: Haskin, D., "Default Route Advertisement in BGP2 and BGP3 Versions of The Border Gateway Protocol", Bolt, Beranek and Newman, Inc., January 1993.
- RFC 1398: Kastenholz, F., "Definitions of Managed Objects for the Ethernet-like Interface Types", FTP Software, Inc., January 1993.
- RFC 1400: Williamson, S., "Transition and Modernization of the Internet Registration Service", Network Solutions, Inc., March 1993.
- RFC 1401: Internet Architecture Board, Lyman Chapin, Chair, "Correspondence between the IAB and DISA on the use of DNS throughout the Internet", IAB, January 1993.
- RFC 1402: Martin, J., "There's Gold in Them Thar Networks! or Searching for Treasure in all the Wrong Places", Ohio State University, January 1993.
- RFC 1403: Varadhan, K., "BGP OSPF Interaction", OARnet, January 1993.
- RFC 1404: Stockman, B., "A Model for Common Operational Statistics", NORDUnet/SUNET, January 1993.
- RFC 1405: Allocchio, C., "Mapping between X.400 (1984/1988) and Mail-11 (DECnet mail)", I.N.F.N. – Italy, January 1993.
- RFC 1406: Baker, F. (Advanced Computer Communications), and J. Watt (Newbridge Networks Corporation), Editors, "Definitions of Managed Objects for the DS1 and E1 Interface Types", January 1993.

- RFC 1407: Cox, T., and K. Tesink, "Definitions of Managed Objects for the DS3/E3 Interface Type", Bell Communications Research, January 1993.
- RFC 1408: Borman, D., Ed., "Telnet Environment Option", Cray Research, Inc., January 1993.
- RFC 1409: Borman, D., Ed., "Telnet Authentication Option", Cray Research, Inc., January 1993.
- RFC 1410: Postel, J., Ed., "IAB Official Protocol Standards", Internet Architecture Board, March 1993.
- RFC 1411: Borman, D., Ed., "Telnet Authentication: Kerberos Version 4", Cray Research, Inc., January 1993.
- RFC 1412: Alagappan, K., "Telnet Authentication: SPX", Digital Equipment Corporation, January 1993.
- RFC 1413: St.Johns, M., "Identification Protocol", ARPA/CSTO, January 1993.
- RFC 1414: St.Johns, M. (ARPA/CSTO), and M. Rose (Dover Beach Consulting, Inc.), "Identification MIB", January 1993.
- RFC 1415: Mindel, J., and R. Slaski, "FTP-FTAM Gateway Specification", Open Networks, Inc., January 1993.
- RFC 1416: Borman, D., Ed., "Telnet Authentication Option", Cray Research, Inc., February 1993.
- RFC 1417: The North American Directory Forum, "NADF Standing Documents: A Brief Overview", February 1993.
- RFC 1418: Rose, M., "SNMP over OSI", Dover Beach Consulting, Inc., March 1993.
- RFC 1419: Minshall, G. (Novell, Inc.), and M. Ritter (Apple Computer, Inc.), "SNMP Over AppleTalk", March 1992.
- RFC 1420: Bostock, S., "SNMP Over IPX", Novell, Inc. March 1993.
- RFC 1421: Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", February 1993.
- RFC 1422: Kent, S., "Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", BBN Communications, February 1993.
- RFC 1423: Balenson, D., "Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers", Trusted Information Systems, February 1992.

- RFC 1424: Kaliski, B., "Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services", RSA Laboratories, February 1993.
- RFC 1425: Klensin, J. (UN University), N. Freed (Innosoft Int'l), M. Rose (Dover Beach Consulting, Inc.), E. Stefferud (Network Management Associates, Inc.), and D. Crocker (The Branch Office), "SMTP Service Extensions", February 1993.
- RFC 1426: Klensin, J. (UN University), N. Freed (Innosoft Int'l), M. Rose (Dover Beach Consulting, Inc.), E. Stefferud (Network Management Associates, Inc.), and D. Crocker (The Branch Office), "SMTP Service Extension for 8bit-MIMEtransport", February 1993.
- RFC 1427: Klensin, J. (UN University), N. Freed (Innosoft Int'l), M. Rose (Dover Beach Consulting, Inc.), E. Stefferud (Network Management Associates, Inc.), and D. Crocker (The Branch Office), "SMTP Service Extension for Message Size Declaration", February 1993.
- RFC 1428: Vaudreuil, G., "Transition of Internet Mail from Just-Send-8 to 8bit-SMTP/MIME", CNRI, February 1993.
- RFC 1429: Thomas, E., "Listserv Distribute Protocol", Swedish University Network, February 1993.
- RFC 1430: Hardcastle-Kille, S. (ISODE-Consortium), E. Huizer (SURFnet bv), V. Cerf (Corporation for National Research Initiatives), R. Hobby (University of California, Davis), and S. Kent (Bolt, Beranek and Newman), "A Strategic Plan for Deploying an Internet X.500 Directory Service", February 1993.
- RFC 1431: Barker, P., "DUA Metrics", University College London, February 1993.
- RFC 1432: Quarterman, J., "Recent Internet Books", Matrix Information and Directory Services, Inc., March 1993.
- RFC 1433: Garrett, J. (AT&T Bell Labs), J. Hagan,(Univ. Penn), and J. Wong (AT&T Bell Labs), "Directed ARP", March 1993.
- RFC 1434: Dixon, R., and D. Kushi, "Data Link Switching: Switch-to-Switch Protocol", IBM, March 1993.
- RFC 1435: Knowles, S., "IESG Advice from Experience with Path MTU Discovery", FTP Software, March 1993.

RFC 1436: Anklesaria, F., M. McCahill, P. Lindner, D. Johnson, D. Torrey, B. Alberti, "The Internet Gopher Protocol (a distributed document search and retrieval protocol)", University of Minnesota, March 1993.

RFC 1439: Linseth, C., "The Uniqueness of Unique Identifiers", University of Minnesota, March 1993.

8.E. VISITORS

- February 11 Discussions were held at ISI with Boesch, Lipton, and Bershad re: PC-ATOMIC. This task involves creation of a 486-based interface to the ATOMIC LAN that will be integrated into a VL-bus equipped 486 motherboard. Support for a version of Mach will be provided. Special attention will be paid to the support of direct transfer between network buffers and user memory.
- February 26 Mike St.Johns visited ISI to discuss Gigabit projects.
- March 9 ATOMIC team discussed ATOMIC and Netstation with Information Presentation Technologies, Inc. at ISI.
- March 19 ATOMIC team gave a presentation on ATOMIC and Netstation to Integrated Device Technology, Inc.

8.F. TRAVEL

Walt Prue attended the Merit/NSFNET Regional Techs meeting in Boulder, CO, January 20-23, 1993.

Eve Schooler attended presentations at the Val Verde Unified School District in Riverside, CA, January 28, 1993.

Peter Will and Greg Finn attended the ARPA High Definition Systems Information Exchange Conference in Arlington, VA, February 1-4, 1993, where Greg Finn gave a presentation on Netstation and network vs. bus inter-subsystem communication.

Bob Felderman gave a presentation on ATOMIC at the University of Virginia and at the CNRI XTP discussions, February 2-8, 1993.

Joe Touch and Walt Prue travelled to ARPA to install ZAPT (formerly ACT) equipment for ZAPT teleconferencing, February 9-14, 1993.

Steve Casner, Greg Finn, and Jeff LaCoss visited the Sun Microsystems HRV research team on February 16, 1993. Greg Finn gave a presentation on Netstation and ATOMIC and Steve Casner discussed HRV workstation collaboration with our conferencing work. Discussions were held on possible future R&D collaboration.

Bob Felderman gave a presentation at NASA Institute for Computer Applications in Science and Engineering, March 12, 1993. Bob Felderman also gave a presentation during the 4th Maryland Workshop on Very High Speed Networks, March 15-17, 1993.

Bob Felderman presented a talk at the ARPA CSTO Software PI meeting on March 18.

Jon Postel, Bob Braden, Eve Schooler, Joyce Reynolds, Steve Casner and Eve Schooler attended the IETF meetings in Columbus Ohio, March 29-April 2.

8.G. SEMINARS

HARRICK M. VIN from UCSD, Multimedia Lab., Dept. of Computer Science and Engineering gave a seminar on "Architectures And Algorithms For Multimedia"

SUGIH JAMIN from USC/ Dept. of Computer Science, gave a talk titled, "An Admission Control Algorithm for Predictive Real-Time Service".

DORAB PATEL from Twin Sun, Inc., discussed "Issues in building a collaborative toolkit".

STUART G. STUBBLEBINE discussed "Protecting the Integrity of Privacy-enhanced Electronic Mail with DES-Based Authentication Codes".

ANNETTE DESCHON presented a short discussion of a proposed experiment to test ISI's security "Tunnel" at ISI, prior to its installation at ARPA. Annette explained the experiment and the anticipated impact, both on Division-7 people and on the rest of the Institute. The objective was to get comments and suggestions, and to discuss any possible drawbacks during the planning phase of this experiment.

9. Publications, papers, and presentations

PUBLICATIONS

Cohen, Danny, Gregory Finn, Robert Felderman, and Annette DeSchon, "ATOMIC: A Low-Cost, Very High-Speed Local Communication Architecture", to appear in the *Proceedings of the International Conference on Parallel Processing*, August 1993.

Schooler, Eve M., "Case Study: Multimedia Conference Control in a Packet-switched Teleconferencing System", to appear in the *Journal of Internetworking*.

Schooler, Eve M., "The Impact of Scaling on a Multimedia Connection Architecture", to appear in the *Journal of Multimedia Systems*.

PAPERS

A memo on "How to do an IETF A/V multicast" was written to help the organizers of this and future IETFs to set up the necessary equipment.

PRESENTATIONS

Presentations on ATOMIC were given to:

- High Definition Workshop (Finn)
- University of Virginia CS Department (Felderman)
- Sun Microsystems (Finn, Casner, LaCoss)
- Information Presentation Technologies, Inc. (Finn, Felderman)
- NASA ICASE Virginia (Felderman)
- 4rd Maryland Workshop on VHSN (Felderman,)
- Gigabit Networks Symposium (Felderman)
- Integrated Device Technology, Inc. (Cohen)